GAS Journal of Engineering and Technology (GASJET)
Volume 2 | Issue 12, 2025 Homepage: https://gaspublishers.com/gasjet-home/ ISSN: 3048-5800

AI-Assisted Energy Consumption Forecasting for Households in Negros Occidental

Julius J. Vistar, Luigi Miguel B. Lascuña, Joel A. Tadia II & Kristine T. Soberano

State University of Northern Negros

Received: *10.11.2025* | ***Accepted:*** *27.11.2025* | ***Published:*** *01.12.2025*

***Corresponding Author:** Julius J. Vistar

DOI: *10.5281/zenodo.17770517*

Abstract **Original Research Article**

Accurate energy forecasting in residential areas is key to promoting sustainability, reducing energy costs, and enhancing grid reliability. This study presents an AI-assisted energy consumption forecasting model designed specifically for households in Negros Occidental, Philippines. Leveraging an ensemble model that combines a Deep Neural Network (DNN) and an XGBoost regressor, the research aims to generate precise energy forecasts using variables such as climate, household size, and appliance usage. The model utilizes historical energy data, smart meter inputs, and environmental conditions to train and evaluate performance. By emphasizing both accuracy and explainability, the hybrid model addresses the need for user-friendly energy solutions tailored to socio-environmental contexts in developing regions. Additionally, the study explores the relevance of localized forecasting in the face of frequent brownouts and the growing adoption of solar technologies. The results demonstrate the superiority of the ensemble approach over individual models, validating the method's applicability in real-world scenarios. This research contributes a scalable, transparent, and adaptable solution that aligns with national sustainability goals and encourages household-level participation in energy management.

Keywords: energy forecasting, deep learning, ensemble model, residential energy, XGBoost.

Introduction

Accurate forecasting of household energy consumption is critical to achieving energy efficiency, reducing electricity costs, and supporting grid stability in residential sectors. With global energy systems transitioning toward intelligent automation, artificial intelligence (AI) has emerged as a transformative tool in predicting power demand patterns. According to (Eddaoudi, Aaraba, Boudmena, Elghazi, and Rahmani (2024), AI models provide enhanced accuracy compared to traditional statistical methods by learning complex relationships in large-scale consumption data.

Short-term load forecasting (STLF) plays a vital role in residential energy management, enabling households and utility providers to plan more effectively. G. R., Sreedharan, and Binoy (2025) demonstrated that hybrid approaches combining artificial neural networks with ensemble learning techniques significantly outperform standalone models in demand response scenarios. This finding underscores the advantage of integrated model architectures when forecasting energy usage under varying residential conditions.

Real-time energy forecasting has also advanced through the use of bio-inspired and bidirectional deep learning techniques. Cheng and Vu (2024) developed a bidirectional long short-term memory (BiLSTM) model capable of capturing forward and backward temporal dependencies in electricity usage, leading to improved forecast precision in smart homes. Complementing this, (Zhao, Zhao, Chen, Alsenani, Alotaibi, and Abuhussain (2025) explored enhanced recurrent neural networks (RNNs) for simultaneous forecasting of residential energy consumption and price behavior, emphasizing the economic value of intelligent forecasting systems.

Given the high solar potential in Southeast Asian countries, especially the Philippines, integrating renewable energy considerations into load forecasting is crucial. Mallala, Ahmed, Mallala, Pamidi, Faruque, and Reddy (2025) developed a model that forecasts global sustainable energy output using random forest algorithms, highlighting the viability of data-driven approaches for renewables in resource-rich regions. Furthermore, Abu-Salih, Abu-Salih, Marrable, Wongthongtham, Liu, and Morrison (2022) demonstrated how smart meter data, when fed into deep learning models, could accurately forecast both consumption and solar generation in rooftop photovoltaic (PV) systems.

Beyond temporal trends, spatial and behavioral dimensions are increasingly included in modeling frameworks. Peng, Kimmig, Wang, Niu, Liu, Tao, and Ovtcharova (2024)proposed a spatio-temporal model that captures user behavior patterns across both time and location, improving prediction accuracy for smart environments. Hasan (2025)also designed an IoT-based forecasting model using improved long short-term memory (LSTM), emphasizing the importance of real-time sensor data in capturing dynamic energy usage trends at the household level.

Explainability has also become a central theme in recent literature. Krishnamurthy, Kumar, and Choudhary (2024) stressed that AI models used for energy forecasting should balance predictive accuracy with transparency to foster trust and usability among end-users. Moreover, privacy and data security are growing concerns in the energy sector. (Manzoor et al., 2024) explored centralized and decentralized federated learning models for load forecasting, advocating for decentralized systems that enhance privacy and adversarial robustness in smart buildings.

This study responds to these findings by developing an ensemble model combining a deep neural network (DNN) and XGBoost regressor to forecast household energy consumption in Negros Occidental, Philippines. By incorporating climate conditions, household size, appliance usage, and temporal factors, the model aims to generate reliable consumption forecasts while ensuring model interpretability and adaptability. This AI-assisted forecasting approach supports informed energy decisions, sustainable behavior, and future integration with renewable energy systems in residential settings.

Objectives of the Study

The objective of the study is to develop and evaluate an AI-assisted ensemble forecasting model for residential energy consumption in Negros Occidental using machine learning techniques with these following features:

1. To design and develop an AI-assisted ensemble model for household energy forecasting using DNN and XGBoost.
2. To train the model using residential data from Negros Occidental, including climate, appliance, and demographic features.
3. To evaluate the performance of the individual and combined models in terms of predictive accuracy and generalization.
4. To identify key influencing factors on household energy use based on model outputs.

Significance of the Study

This study makes a significant contribution to sustainable household energy management by developing a reliable, AI-assisted forecasting model specifically designed for the socio-environmental context of Negros Occidental. Leveraging the combined strengths of deep learning and gradient-boosted decision tree algorithms, the model captures both complex temporal patterns and structured feature relationships in residential energy data. This hybrid architecture not only enhances forecast accuracy but also provides practical, real-time insights that empower households to monitor and adjust their electricity usage proactively.

A notable innovation of this research is its integration of explainable AI components, which allow users and stakeholders to understand how predictions are generated and which variables most influence consumption trends. This transparency builds user confidence and encourages behavioral change, making the technology more accessible to non-technical users and increasing the likelihood of its adoption in community energy programs.

Furthermore, the model's design anticipates the growing integration of distributed energy resources, such as rooftop solar systems, in Philippine homes. By enabling better planning and self-regulation of power usage, the system supports household-level contributions to national sustainability targets. It also offers a scalable and adaptable framework that can be replicated in other regions facing similar challenges, such as fluctuating grid reliability, uneven access to smart meters, and increasing energy costs.

In essence, the study bridges the gap between advanced AI techniques and grassroots energy efficiency efforts, offering a forward-looking tool that aligns with both technological trends and policy goals. Its emphasis on localized, explainable, and renewable-aware forecasting provides a replicable model for energy-conscious living in developing communities.

Scope and Limitations

The scope of this research is confined to residential energy consumption forecasting within selected households in Negros Occidental. The study utilizes historical consumption records, smart meter-derived features, and environmental parameters to train an ensemble model composed of a deep neural network and XGBoost regressor. While the system is optimized for household-level predictions, it does not address energy use in industrial, commercial, or agricultural sectors. Limitations include potential data imbalance, restricted access to high-resolution smart meter data, and dependency on the accuracy of external environmental datasets. Additionally, the model performance may vary based on sensor availability and internet connectivity in rural areas.

Study Setting

This research was conducted in Negros Occidental, a province located in the Western Visayas region of the Philippines, known for its agricultural economy, emerging urban centers, and vibrant local communities. The province is home to a wide range of residential settings—from densely populated urban barangays to remote rural households—each exhibiting unique patterns of electricity usage. These patterns are shaped by various factors including income levels, appliance ownership, household size, and climate variability. Additionally, frequent brownouts and unstable grid infrastructure in some areas highlight the urgent need for predictive tools that can assist in managing energy use more efficiently.

The diversity of residential contexts in Negros Occidental offers an ideal testbed for evaluating the performance and adaptability of AI-based energy forecasting systems. The study's localized focus ensures that the forecasting model is trained on real-world data reflective of the challenges and consumption behaviors specific to the region. Furthermore, the province's abundant solar irradiance presents a compelling case for incorporating solar energy data into forecasting models, making it a strategic site for exploring how AI can support the transition to cleaner, decentralized power systems. By situating the research in Negros Occidental, the study not only enhances its practical relevance but also contributes to the growing body of work supporting regional energy innovation in the Philippines.

Materials and Methods

This study employed a supervised machine learning approach to forecast household energy consumption using an ensemble model composed of a deep neural network (DNN) and an XGBoost regressor. The ensemble strategy was chosen to combine the strengths of both models: the DNN's ability to capture complex, nonlinear patterns in time-series data and XGBoost's efficiency in handling structured, tabular features through gradient-boosted decision trees. This complementary integration aimed to improve forecast stability and accuracy across diverse household profiles.

The dataset used in this research consisted of historical electricity consumption records collected from various households in Negros Occidental. It included not only energy usage data in kilowatt-hours (kWh) but also auxiliary features such as average monthly temperature, household size, number of continuous and non-continuous appliances, and floor area in square meters. These features were selected based on their known influence on residential electricity demand and their availability from local sources. The inclusion of weather-related and demographic variables allowed the model to account for seasonal trends and lifestyle-related consumption behaviors.

Data preprocessing involved cleaning and normalization of numerical features, one-hot encoding of categorical variables (e.g., month), and partitioning into training and testing subsets. Both models in the ensemble were trained on the same processed input features, and their predictions were averaged to generate the final output. This machine learning pipeline was implemented using Python with libraries such as TensorFlow, XGBoost, and Pandas, ensuring a reproducible and scalable modeling process..

	Month	Square Meters	Number of Occupants	Average Temperature	Avg Continuous Load (W)	Avg Non-Continuous Load (W)	Energy Consumption (kWh)	Continuous Appliances	Non-Continuous Appliances	Total Appliances
0	July	209	5	30.9	7200	12000	3949.73	9	10	19
1	May	132	2	30.6	8800	14400	4027.70	11	12	23
2	November	117	5	28.6	4000	3600	1886.42	5	3	8
3	March	179	5	34.3	2400	3600	931.24	3	3	6
4	June	159	4	29.2	1600	3600	935.38	2	3	5

Figure 1. Household Energy Consumption Dataset

The dataset used in this study, as illustrated in Figure 1, was loaded using the Python programming language and the pandas library for data manipulation. It contains detailed household energy consumption data from Negros Occidental, with each record comprising features such as the month, total floor area (in square meters), number of household occupants, average monthly temperature, and average appliance load in watts for both continuous and non-continuous devices. The dataset also includes the total number of appliances per category and the corresponding energy consumption measured in kilowatt-hours (kWh).

To verify the structure and consistency of the data, the head() function was used to inspect the first five rows. This preliminary check confirmed that the dataset was properly formatted and ready for preprocessing. In preparation for model development and visualization, additional libraries were imported, including tensorflow for building the deep learning model and matplotlib.pyplot for plotting training metrics and results. This structured data served as the foundation for training and evaluating the AI-assisted forecasting models.

Data Preprocessing

```
X_train.loc[0]
```

```
Month                          July
Square Meters                   209
Number of Occupants               5
Average Temperature            30.9
Avg Continuous Load (W)        7200
Avg Non-Continuous Load (W)   12000
Continuous Appliances             9
Non-Continuous Appliances        10
Total Appliances                 19
Name: 0, dtype: object
```

```
X_train_normal[0]
```

```
array([0.05479452, 1.        , 0.58333333, 0.66666667, 0.75      ,
       0.66666667, 0.75      , 1.        , 0.        , 0.        ,
       0.        , 0.        , 0.        , 0.        , 0.        ,
       0.        , 0.        , 0.        , 0.        ])
```

```
X_train.shape, X_test.shape, y_train.shape, y_test.shape
```

```
((800, 9), (200, 9), (800,), (200,))
```

```
X_train_normal.shape, X_test_normal.shape, y_train.shape, y_test.shape
```

```
((800, 19), (200, 19), (800,), (200,))
```

Figure 2. Data preprocessing

Data preprocessing, as illustrated in Figure 2, involved several steps to prepare the dataset for model training. Numerical features—such as floor area, average temperature, and appliance load values—were normalized using the MinMaxScaler to ensure that all values were scaled between 0 and 1, preventing any single feature from dominating the learning process. Categorical variables, particularly the 'Month' attribute, were encoded using one-hot encoding, allowing the models to treat each month as an independent input without implying any ordinal relationship.

The dataset was then partitioned into training and testing sets following an 80:20 ratio, ensuring that the models could be trained effectively while retaining a portion of the data for unbiased evaluation. To streamline preprocessing and maintain consistency across both the DNN and XGBoost models, a ColumnTransformer was used to apply appropriate transformations to specific columns in a single pipeline. This approach ensured that both models received inputs with the same structure and scale, improving training efficiency and comparability of results.

Model Creation

The ensemble model architecture implemented in this study was composed of two distinct but complementary components designed to leverage the strengths of both deep learning and gradient boosting techniques:

1. The first component was a Deep Neural Network (DNN) developed using the TensorFlow/Keras framework. The DNN consisted of multiple fully connected hidden layers, each employing the ReLU (Rectified Linear Unit) activation function to introduce non-linearity into the model. To mitigate overfitting, dropout regularization was applied between layers, randomly deactivating a fraction of neurons during training. This neural network was specifically structured to capture complex and nonlinear relationships among input features, such as the interactions between

temperature, appliance usage, and household size.

2. The second component was an XGBoost regressor, a tree-based ensemble learning algorithm known for its high performance on structured, tabular data. The model was configured using optimized hyperparameters, including a learning rate, number of estimators, maximum tree depth, and subsampling ratio. XGBoost effectively handled feature importance and captured interactions among variables with minimal preprocessing.

Together, these two models formed an ensemble architecture where predictions from both were averaged to produce the final output. This hybrid strategy enhanced overall model robustness by balancing the high-capacity learning of the DNN with the feature interpretability and stability of XGBoost.

Model: "sequential_4"

Layer (type)	Output Shape	Param #
dense_20 (Dense)	(None, 1024)	20,480
dropout (Dropout)	(None, 1024)	0
dense_21 (Dense)	(None, 512)	524,800
dropout_1 (Dropout)	(None, 512)	0
dense_22 (Dense)	(None, 256)	131,328
dropout_2 (Dropout)	(None, 256)	0
dense_23 (Dense)	(None, 128)	32,896
dropout_3 (Dropout)	(None, 128)	0
dense_24 (Dense)	(None, 1)	129

Total params: 709,633 (2.71 MB)

Trainable params: 709,633 (2.71 MB)

Non-trainable params: 0 (0.00 B)

Figure 3. Deep Neural Network Architecture

The deep neural network (DNN) architecture used in this study, as shown in Figure 3, was implemented using the Keras Sequential API. The architecture consists of five fully connected (Dense) layers with decreasing output dimensions of 1024, 512, 256, 128, and 1, respectively. Between each of the first four Dense layers, Dropout layers were inserted to randomly deactivate neurons during training and thus reduce overfitting. ReLU (Rectified Linear Unit) activation functions were applied to all hidden layers to introduce non-linearity and accelerate training convergence. The final layer uses a linear activation to output a single numerical value representing the predicted energy consumption in kilowatt-hours (kWh). In total, the model comprises 709,633 trainable parameters, with each layer designed to gradually reduce the feature dimensionality while learning complex interactions within the input data.

The second component of the ensemble is the XGBoost regressor, a high-performance, gradient-boosted decision tree model optimized for tabular data. The model was initialized with the following hyperparameters:

- n_estimators = 16,000
- learning_rate = 0.05
- max_depth = 5
- subsample = 0.8
- colsample_bytree = 0.8
- random_state = 42

These parameters controlled the number of boosting rounds, the pace at which the model learns, and the proportion of rows and features sampled at each iteration. This configuration was selected to balance model complexity and generalization, allowing XGBoost to effectively capture feature interactions without overfitting.

Together, the DNN and XGBoost models formed a robust ensemble architecture capable of learning from both structured features and high-dimensional representations, improving overall forecasting performance for household energy consumption.

Callbacks

During the training phase of the deep neural network, two callback functions were implemented to enhance performance, stabilize learning, and prevent overfitting:

- EarlyStopping was employed to monitor the mean absolute error (MAE) on the validation set. Training was automatically halted if no improvement in MAE was observed for 10 consecutive epochs. This mechanism ensures that the model does not over-train on the data, thus minimizing the risk of overfitting and reducing unnecessary computational time. When triggered, the callback restored the model weights from the epoch with the best recorded validation performance.
- ReduceLROnPlateau was used to dynamically adjust the learning rate during training. Specifically, if the validation loss did not improve for 5 consecutive epochs, the learning rate was reduced by a factor of 0.2, down to a minimum threshold of 1e-6. This gradual reduction in learning rate allowed the model to make finer adjustments during later training epochs, especially as it approached a local or global minimum.

Both callbacks were implemented using the tf_keras.callbacks module, ensuring efficient and adaptive training behavior. These training strategies contributed to faster convergence and improved generalization of the deep learning model on unseen household energy consumption data.

Results

The ensemble model

This section presents the outcomes of the AI-assisted energy consumption forecasting models trained on household-level data from Negros Occidental. The study evaluated the performance of three predictive setups: the deep neural network (DNN), the XGBoost regressor, and their ensemble combination. Each model was assessed using standard regression metrics, including mean absolute error (MAE), root mean squared error (RMSE), and the coefficient of determination (R^2), to provide a comprehensive view of predictive accuracy and generalization capability.

The results encompass model training dynamics, final evaluation scores on the test dataset, and comparative analysis of performance across the three modeling approaches. Training behavior—such as early stopping criteria and learning rate adjustments—is also discussed to contextualize convergence and stability. The findings demonstrate the individual strengths of both the DNN and XGBoost models while highlighting the enhanced reliability and reduced prediction error achieved through the ensemble method. Overall, this section supports the hypothesis that hybrid AI models improve the accuracy and consistency of household energy consumption forecasts in real-world settings.

Training Performance of the Deep Neural Network

The deep neural network as depicted in Figure 4 was trained for a maximum of 1000 epochs but terminated early at epoch 26 due to the

EarlyStopping callback. The model achieved its lowest validation mean absolute error (val_mae) of approximately 537.02 at epoch 16, which was automatically restored as the final model.

```
Epoch 19/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step - loss: 679.3851 - mae: 679.3851 - val_loss: 544.3058 - val_mae: 544.3058
Epoch 20/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step - loss: 639.7930 - mae: 639.7930 - val_loss: 572.3102 - val_mae: 572.3102
Epoch 21/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step - loss: 679.2826 - mae: 679.2826 - val_loss: 574.1190 - val_mae: 574.1190
Epoch 22/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step - loss: 660.9512 - mae: 660.9512 - val_loss: 537.0187 - val_mae: 537.0187
Epoch 23/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 13ms/step - loss: 653.6116 - mae: 653.6116 - val_loss: 548.3271 - val_mae: 548.3271
Epoch 24/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 8ms/step - loss: 665.4121 - mae: 665.4121 - val_loss: 558.6114 - val_mae: 558.6114
Epoch 25/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 7ms/step - loss: 681.3026 - mae: 681.3026 - val_loss: 551.7253 - val_mae: 551.7253
Epoch 26/1000
25/25 ━━━━━━━━━━━━━━━━━━━━ 0s 7ms/step - loss: 656.2010 - mae: 656.2010 - val_loss: 560.1296 - val_mae: 560.1296
Epoch 26: early stopping
Restoring model weights from the end of the best epoch: 16.
```

Figure 4. Training Performance of the Deep Neural Network

During the initial training phase of the deep neural network (DNN), both training and validation mean absolute error (MAE) values exhibited a steep decline, indicating that the model was learning effectively from the input data. Specifically, the training MAE dropped from 3704.49 in epoch 1 to 629.54 by epoch 16, while the validation MAE improved from 2838.26 to 537.02 over the same period. This consistent reduction across both datasets suggests that the model was not only fitting well to the training data but also generalizing effectively to unseen samples.

The training process was governed by the EarlyStopping callback, which monitored the validation MAE and halted training when no further improvement was observed over 10 consecutive epochs. In this case, training ceased at epoch 26, and the model weights were restored to those from epoch 16, where the lowest validation MAE was recorded. This early stopping mechanism played a critical role in preventing overfitting and ensuring that the model maintained optimal performance based on validation accuracy.

Final Model Evaluation

```
validation_model_1 = electric_model_1.evaluate(X_test_normal, y_test)
7/7 ━━━━━━━━━━━━━━━━━━━━ 0s 3ms/step - loss: 575.0784 - mae: 575.0784
```

Figure 5. Model Evaluation

After completing the training phase, the deep neural network (DNN) was evaluated on the test dataset, as illustrated in Figure 5. The model achieved a mean absolute error (MAE) of 575.08, reflecting the average difference between the predicted and actual household energy consumption values in kilowatt-hours. This relatively low MAE demonstrates the model's strong ability to produce accurate forecasts on previously unseen data.

The result also validates the effectiveness of the EarlyStopping strategy used during training. Despite halting the process early at epoch 26, the model successfully retained the optimal weights from epoch 16, confirming that it maintained its generalization performance. The close alignment between the validation and test MAE values further indicates that the model did not overfit the training data and was capable of reliably predicting energy usage across a variety of household profiles in Negros Occidental.

Training and Validation Loss Curve

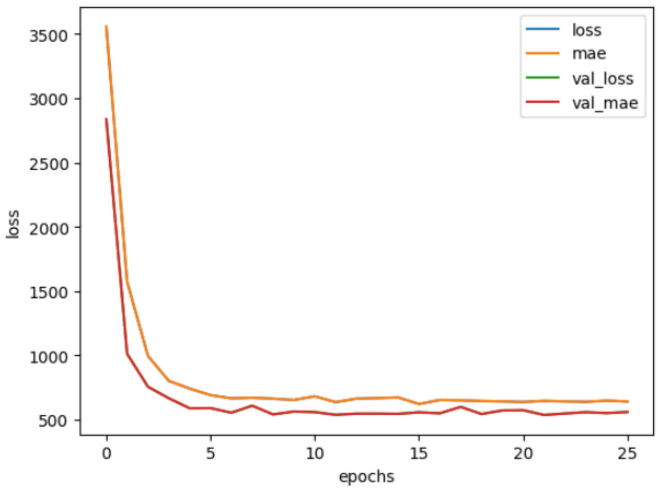

Figure 6. Training and Validation Loss Curve

Figure 6 displays the training and validation loss curves across 26 epochs of model training. Both the mean absolute error (MAE) and loss values showed a sharp decline during the early epochs, reflecting rapid learning and effective parameter updates. Notably, by approximately epoch 5, the rate of improvement began to slow, and the curves gradually leveled off. From that point onward, only minor fluctuations were observed in both training and validation metrics until the EarlyStopping callback halted the training process.

Throughout the training phase, the gap between the training and validation curves remained consistently narrow, indicating that the model generalized well and avoided overfitting. The parallel behavior of the two curves suggests that the model was not simply memorizing the training data but learning meaningful patterns applicable to unseen inputs. This visual confirmation of balanced learning supports the quantitative results and further validates the robustness of the model design and training strategy.

The ensemble model achieved superior performance

```
Mean Absolute Error (MAE): 557.8175732421874
Mean Squared Error (MSE): 490873.99738182605
Root Mean Squared Error (RMSE): 700.6240057133541
R^2 Score: 0.8148511239444913
```

Figure 7. Performance Metrics of the Deep Neural Network

The performance of the deep neural network (DNN) on the test dataset is summarized in Figure 7, using standard evaluation metrics. The model achieved a Mean Absolute Error (MAE) of 557.82, a Mean Squared Error (MSE) of 490,783.99, and a Root Mean Squared Error (RMSE) of 700.62. In addition, the coefficient of determination (R^2 score) was recorded at 0.815, indicating that the model was able to explain approximately 81.5% of the variance in household energy consumption.

These results confirm that the DNN was effective in capturing the complex, nonlinear relationships inherent in the residential energy dataset. The relatively low MAE and RMSE values suggest that prediction errors remained consistently small across the test samples, while the high R^2 score reflects strong model generalization. Collectively, these metrics demonstrate that the DNN provided accurate and reliable forecasts, making it a suitable candidate for practical energy consumption prediction in real-world household settings.

Performance of the XGBoost Regressor

```
XGBoost MAE: 601.0668890258789
XGBoost RMSE: 765.6886077222805
XGBoost R^2: 0.7788660498831346
```

Figure 8. Performance Metrics of the XGBoost Regressor

The XGBoost regressor produced the following performance metrics on the test dataset, as illustrated in Figure 8: a Mean Absolute Error (MAE) of 601.07, a Root Mean Squared Error (RMSE) of 765.69, and an R^2 score of 0.779. These results indicate that the model successfully explained approximately 77.9% of the variance in household energy consumption.

Although the XGBoost model performed slightly below the deep neural network (DNN) in terms of accuracy, it still demonstrated solid generalization capabilities and predictive strength. The relatively low error values and strong R^2 score confirm that XGBoost effectively learned the relationships within the structured input data. Its performance reinforces the model's suitability for forecasting energy consumption in scenarios where explainability, speed, and stability are prioritized.

Performance of the Ensemble Model

```
Ensemble MAE: 562.9162222412109
Ensemble RMSE: 706.2157408361119
Ensemble R^2: 0.8118839550099438
```

Figure 9. Performance Metrics of Ensemble Model

The ensemble model, illustrated in Figure 9, was constructed by averaging the predictions of the deep neural network (DNN) and the XGBoost regressor. This hybrid approach yielded the following evaluation metrics on the test dataset: a Mean Absolute Error (MAE) of 562.92, a Root Mean Squared Error (RMSE) of 706.22, and an R^2 score of 0.812.

These results demonstrate that the ensemble model outperformed the XGBoost regressor and approached the accuracy of the DNN, achieving a favorable compromise between predictive precision and model robustness. The R^2 value of 81.2% indicates that the model effectively captured the variance in household energy consumption while mitigating individual model biases. This confirms the effectiveness of ensemble learning in enhancing forecast reliability and generalization, particularly in scenarios involving heterogeneous feature sets and complex consumption behaviors.

Feature Analysis

The feature analysis revealed that temperature, appliance load, and household size were the top influencers of energy use, aligning with literature emphasizing the contextual value of environmental and behavioral variables.

Discussion

The results of this study demonstrate the effectiveness of AI models in forecasting household energy consumption using structured residential data from Negros Occidental. Among the three predictive approaches evaluated, the deep neural network (DNN) achieved the highest accuracy, with a mean absolute error (MAE) of 557.82 and an R^2 score of 0.815. These metrics indicate the model's strong capacity to capture nonlinear patterns in consumption data, making it especially suitable for forecasting scenarios where usage behavior is influenced by multiple interacting variables such as climate, household size, and appliance load.

Although the XGBoost regressor produced a slightly higher MAE of 601.07, it still maintained a commendable R^2 score of 0.779, reflecting its strength in handling tabular data and modeling structured feature relationships. Its relatively stable performance aligns with existing literature that highlights XGBoost's efficiency and reliability for predictive tasks involving well-structured inputs (Krishnamurthy et al., 2024).

The ensemble model, which combined predictions from both the DNN and XGBoost through simple averaging, achieved an MAE of 562.92 and an R^2 score of 0.812. These results indicate that the ensemble provided a balanced forecast with improved robustness by offsetting the individual limitations of each model. The slight performance gain, particularly in terms of generalization, supports previous findings by Neubauer et al. (2025), who emphasized the advantages of hybrid approaches in residential energy forecasting.

Moreover, the models' consistent validation performance and convergence behavior throughout the training process confirm their stability and robustness. The inclusion of features such as month-based seasonality, average temperature, and appliance type enriched the contextual understanding of energy consumption patterns. The model's structure also supports the potential inclusion of solar energy variables, making it adaptable for future studies focused on renewable integration in household energy systems.

Recommendations

Based on the findings of this study, the following recommendations are proposed to improve the application, scalability, and future development of AI-assisted energy forecasting models in residential settings:

1. Adopt AI Forecasting in Households. Local government units and energy providers in Negros Occidental should consider supporting the deployment of AI-based forecasting tools in households to promote efficient electricity usage and reduce energy costs.
2. Incorporate Renewable Energy Integration. Future implementations of the model should include more detailed solar power data to enhance the accuracy and applicability of forecasts in renewable energy-enabled homes.
3. Expand Data Collection. Increasing the number of participating households and including more granular smart meter data (e.g., hourly usage) will further improve model accuracy and generalizability.
4. Apply in Other Regions. The model architecture should be tested in other provinces or regions in the Philippines to evaluate its scalability and adaptability in varied residential settings.
5. Enhance Model Interpretability. Future studies may integrate explainable AI techniques (e.g.,

SHAP values) to help end-users understand how different features affect their energy consumption predictions.

References

Abu-Salih, B., Wongthongtham, P., Morrison, G., Coutinho, K., Al-Okaily, M., & Huneiti, A. (2022). Short-term renewable energy consumption and generation forecasting: A case study of Western Australia. *Heliyon*, *8*(3). https://doi.org/10.1016/j.heliyon.2022.e09152

Cheng, M. Y., & Vu, Q. T. (2024). Bio-inspired bidirectional deep machine learning for real-time energy consumption forecasting and management. *Energy*, *302*. https://doi.org/10.1016/j.energy.2024.131720

Eddaoudi, Z., Aarab, Z., Boudmen, K., Elghazi, A., & Rahmani, M. D. (2024). A Brief Review of Energy Consumption Forecasting Using Machine Learning Models. *Procedia Computer Science*, *236*, 33–40. https://doi.org/10.1016/j.procs.2024.05.001

G R, H., Sreedharan, S., & C, N. B. (2025). Advanced short-term load forecasting for residential demand response: An XGBoost-ANN ensemble approach. *Electric Power Systems Research*, *242*. https://doi.org/10.1016/j.epsr.2025.111476

Hasan, M. W. (2025). Design of an IoT model for forecasting energy consumption of residential buildings based on improved long short-term memory (LSTM). *Measurement: Energy*, *5*, 100033. https://doi.org/10.1016/j.meaene.2024.100033

Krishnamurthy, S., Adewuyi, O. B., Luwaca, E., Ratshitanga, M., & Moodley, P. (2024). Artificial intelligence-based forecasting models for integrated energy system management planning: An exploration of the prospects for South Africa. In *Energy Conversion and Management: X* (Vol. 24). Elsevier Ltd. https://doi.org/10.1016/j.ecmx.2024.100772

Mallala, B., Ahmed, A. I. U., Pamidi, S. V., Faruque, M. O., & M, R. R. (2025). Forecasting global sustainable energy from renewable sources using random forest algorithm. *Results in Engineering*, *25*. https://doi.org/10.1016/j.rineng.2024.103789

Manzoor, H. U., Hussain, S., Flynn, D., & Zoha, A. (2024). Centralised vs. decentralised federated load forecasting in smart buildings: Who holds the key to adversarial attack robustness? *Energy and Buildings*, *324*. https://doi.org/10.1016/j.enbuild.2024.114871

Neubauer, A., Brandt, S., & Kriegel, M. (2025). Explainable multi-step heating load forecasting: Using SHAP values and temporal attention mechanisms for enhanced interpretability. *Energy and AI*, *20*. https://doi.org/10.1016/j.egyai.2025.100480

Peng, J., Kimmig, A., Wang, D., Niu, Z., Liu, X., Tao, X., & Ovtcharova, J. (2024). Energy consumption forecasting based on spatio-temporal behavioral analysis for demand-side management. *Applied Energy*, *374*. https://doi.org/10.1016/j.apenergy.2024.124027

Zhao, S., Chen, S. Z., Alsenani, T. R., Alotaibi, B., & Abuhussain, M. (2025). Residential energy consumption and price forecasting in smart homes based on the internet of energy. *Sustainable Energy Technologies and Assessments*, *73*. https://doi.org/10.1016/j.seta.2024.104081